

A 3D Optical Flow Approach to Addition of Deformable PET Volumes

Gregory J. Klein and Ronald H. Huesman
Center for Functional Imaging,
Lawrence Berkeley National Laboratory
University of California
1 Cyclotron Road, MS 55-121
Berkeley, CA 94720
GJKlein@lbl.gov, RHHuesman@lbl.gov

Abstract

A method for combining datasets from gated cardiac PET acquisitions is described. Optical flow techniques are used to accurately model non-rigid motion present during the cardiac cycle so that a one-to-one mapping is found between each voxel of two gated volumes. Using this mapping, images can be combined to produce a composite dataset with improved statistics and reduced motion-induced blur. Like recent past efforts in deformable motion, an image similarity measure is combined with elastic constraints to obtain a valid flow field. Additionally, because of the noisy characteristics of individual reconstructed volumes in a gated PET study, 4D and multiscale techniques were used in this paper to obtain more accurate motion estimates. Results using data from gated cardiac studies on canine and human subjects are presented.

1. INTRODUCTION

Gated acquisition of Positron Emission Tomography (PET) data is a common method for compensation of motion due to the cardiac cycle in imaging of the heart. In this process, tomographic data are collected while monitoring electrocardiograph (ECG) signals. Depending upon the time since the last cardiac R-wave, the data are directed into different histograms. Data from multiple heart beats are accumulated into each histogram. By choosing time segments short enough, the motion of the heart can be effectively frozen as the data are independently reconstructed into separate 3D image volumes. An unfortunate side-effect of this gating is that the statistical quality of each gated image suffers as the total number of PET events is distributed over a number of different

images. One way to improve statistics in a gated study is simply to increase the acquisition time of each histogram to the acquisition time normally used in a static (i.e. non-gated) study. However, this tactic is usually infeasible because of short-lived imaging radiotracers or because of patient's ability to remain stationary in the scanner for this extended time period. An alternate approach for improving image statistics is to sum the gated reconstructions, but to do so without introducing motion-induced blur, a correspondence must first be found between the voxels in each dataset. Only after this one-to-one mapping is obtained can these corresponding voxels be added.

For two volumetric datasets representing an object undergoing a simple rigid-body motion, the correspondence problem amounts to finding the six parameters of translation and rotation which best align two gated images, or gates, as they will be called in this paper. This is the type of motion model which has frequently been used for registering different 3D datasets of the brain [1],[2]. Motion of the heart throughout the cardiac cycle, however, is not rigid. The walls of the heart move as an elastic body during the cardiac contraction. Therefore, rather than a six parameter motion model relating the position of a voxel in one dataset with its corresponding voxel in the other, a dense 3D vector map at the same resolution of the dataset is required to describe the deformations.

This paper describes how 3D optical flow techniques were used to obtain the non-rigid motion estimates between gates in cardiac ^{18}F -fluorodeoxyglucose (FDG) PET studies obtained from canine and human subjects. Because individual gates can be quite noisy, a multi-scale approach was used to insure robust results. Once motion estimates are obtained, the individual gates are deformed to match a reference time, typically chosen to be diastole. Because the composite dataset represents the tomographic events from the entire acquisition, the resulting image sta-

tistics are improved, yet motion induced blur is reduced because of the non-rigid motion compensation.

2. BACKGROUND

Optical flow techniques have historically been used to estimate pixel correspondences between 2D images obtained from two different viewpoints or at two different times. The vector field describing the mapping between corresponding pixels in two images is called a flow field. Classically, researchers have relied on two types of constraints to determine the flow field: first an image matching constraint, and second a smoothness constraint on the resulting flow field. For example, Horn and Schunck [3] make the assumption that for small changes in camera position, the image intensity will be approximately conserved in the two images for pixels corresponding to the same object in space [3]. They combine this constraint with a global motion vector smoothness constraint to obtain a least squares solution of the motion field. Fitzpatrick [4] generalized these ideas to suggest that one-to-one geometric image transformations could be found describing deformable motion for 3D density images, like those typically produced by modern PET, computed tomography (CT), or magnetic resonance imaging (MRI) devices. Song and Leahy [5] expanded these ideas and the optical flow formulation of Horn [3] to produce an algorithm for estimating the 3D transformation between frames in sequence of ultrafast CT cardiac images. Image similarity measures are based on the assumption of conservation of matter, analogous to Horn’s pixel brightness conservation constraint. Like Horn’s algorithm, Song’s 3D implementation imposed a smoothness measure on the flow field to arrive at a unique solution to the problem. In a later work [6], Song added a divergence-free constraint to improve on the accuracy of the flow field. In a related work based on optical flow, Zhou, *et al.* [7] implemented a multiresolution approach using an image matching criterion and non-linear regression which would allow solutions for datasets with large interframe displacements.

A parallel body of work on deformable models has developed as a result of efforts to match brain datasets to a generalized brain atlas. These algorithms invoke a similar image matching criterion as the optical flow techniques, but replace the flow field smoothness and divergence free criterion with more realistic elastic material models. Bajcsy and Kovacic [8] were of the first to propose an implementation, where they used a linear elastic material model given by:

$$\mu \nabla^2 u_i + (\lambda + \mu) \frac{\partial \theta}{\partial x_i} + F_i = 0, (i = 1, 2, 3)$$

Here, the displacement field is defined in their notation as

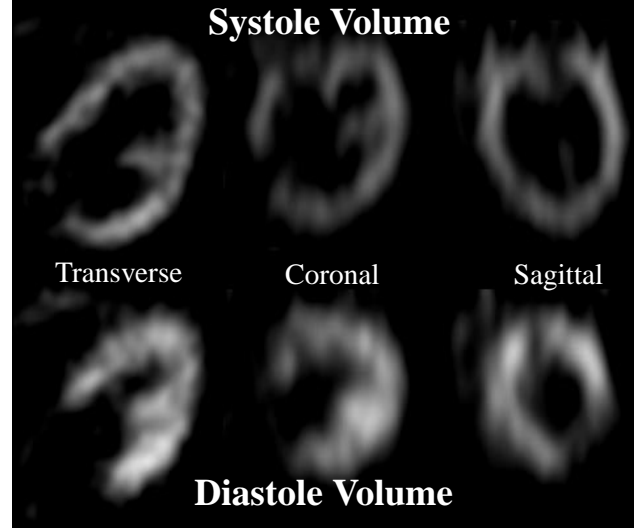


Figure 1. Three orthogonal slices of a gated human cardiac dataset. Top three images show heart during diastole, bottom three show the heart during systole.

$$\mathbf{u} = (u_1, u_2, u_3), \theta = \frac{du_1}{dx_1} + \frac{du_2}{dx_2} + \frac{du_3}{dx_3} = \nabla \cdot \mathbf{u}(x), \lambda \text{ and } \mu$$

are elasticity constants, and F_i is the “body force” derived from the image similarity measure which drives the deformation. Christensen, *et al.* [9],[10] substituted the linear elastic model with a viscous fluid model to arrive at an algorithm more capable of tracking large deformations, however, the added computational complexities are considerable. Other researchers [11],[12], have combined surface-based correspondences with density information to obtain a global deformation which best brings two density datasets into registration.

As seen in Figure 1, the primary feature in FDG PET images of the heart is the left ventricle, which represents the tissue with the greatest radiotracer uptake once the tracer clears the blood pool. Adjacent tissue, such as the atria, right ventricle, blood pool and lung appear primarily as a noisy background. It can be seen from the figure that the movement of corresponding points between diastole and systole can be on the order of several voxels. Because the entire left ventricle remains within the field of view in all gates, it is a reasonable assumption that voxel intensity is conserved. Likewise, because the ventricle is a continuous elastic body, it is reasonable to impose a smoothness criterion for the flow field describing the ventricular motion.

3. DATA ACQUISITION

Cardiac PET data were acquired using the CTI/Sie-

mens ECAT EXACT HR tomograph which we have modified in collaboration with CTI, Inc. for doubly-gated acquisitions. Details of this acquisition system, which allows gating of both the respiratory and the cardiac cycles, may be found in [13]. For the emission study, the subject is first injected with the radiotracer, ^{18}F -fluorodeoxyglucose. After allowing time for the isotope to clear the blood pool, emission data are collected for 20 minutes, distributing the data into 16 different gates. For the canine study, the cardiac cycle was divided into 8 25-msec segments beginning at the peak of the R-wave, and the respiratory cycle was divided into 2 segments, one representing end expiration, the other representing inspiration. In the human study, the cardiac cycle was divided into 11 100-msec segments, and the respiratory cycle was divided into 5 segments. Because of memory limitations in the tomograph hardware, cardiac gating was carried out during only one of the respiratory segments. Acquisition time for the human study was 30 minutes. Prior to the administration of the isotope in both studies, a 20 minute ungated transmission dataset was acquired to correct for the effects of attenuation. Each gate was separately reconstructed into $128 \times 128 \times 47$ voxel volume (voxel size $1.0 \times 1.0 \times 3.1$ mm) using standard filtered backprojection techniques.

4. MOTION ESTIMATION

We use a similar approach to Zhou, *et al.* [7], where the image matching constraint is a function of the difference between a reference volume and the deformed volume and where the solution is regularized by imposing smoothness and incompressibility constraints. In addition, our formulation makes use of 4D information by adding *a priori* motion information from adjacent time frames. The motion estimation framework is described as follows. Define two 3D density fields, $f_1(\mathbf{r})$, and $f_2(\mathbf{r})$, $\mathbf{r} = (x, y, z)$, in a discrete domain,

$$\mathbf{r} \in \Omega = \{[1, N_x], [1, N_y], [1, N_z]\},$$

where N_x, N_y, N_z are the dimensions of the image volume. A motion field is defined as,

$$\mathbf{m}(x, y, z) = (u(x, y, z), v(x, y, z), w(x, y, z))$$

and the deformed volume of f_1 is,

$$\hat{f}_1(\mathbf{r}) = f_1(\mathbf{r} - \mathbf{m}).$$

Assuming f_1 and f_2 are perfect measurements of a conserved medium, the goal of the motion estimation technique is to find the motion field such that

$$\hat{f}_1(\mathbf{r}) = f_2(\mathbf{r}).$$

Since the motion field defines a one-to-one mapping, we can also define the deformed volume of f_2 as

$$\hat{f}_2(\mathbf{r}) = f_2(\mathbf{r} - \mathbf{m}^{-1})$$

With these equations, we wish to calculate the motion field such that the following error measures over each voxel are minimized:

image matching:

$$e_I = \lambda_I(f_2(\mathbf{r}) - \hat{f}_1(\mathbf{r}))^2 + \lambda_I(f_1(\mathbf{r}) - \hat{f}_2(\mathbf{r}))^2$$

smoothness:

$$e_S = \lambda_S(u_x^2 + u_y^2 + u_z^2 + v_x^2 + v_y^2 + v_z^2 + w_x^2 + w_y^2 + w_z^2)$$

incompressibility:

$$e_D = \lambda_D(u_x + v_y + w_z)^2$$

and consistency with a prediction field:

$$e_P = \lambda_P((u - \underline{u})^2 + (v - \underline{v})^2 + (w - \underline{w})^2)$$

where we use the notation, $u_x = \frac{du}{dx}$ and where

$\lambda_I, \lambda_S, \lambda_D, \lambda_P$ are scalars used to weight the different error terms.

A least squares solution to the weighted error terms over all voxels is found via successive linear approximations of $\hat{f}_1(\mathbf{r})$ and $\hat{f}_2(\mathbf{r})$. Assuming the true motion field is \mathbf{m} , and the current estimate of the field is $\tilde{\mathbf{m}}$, then a Taylor series approximation for $\hat{f}_1(\mathbf{r})$ can be defined in terms of $\delta\mathbf{m} = \tilde{\mathbf{m}} - \mathbf{m}$ as:

$$\hat{f}_1(\mathbf{r}) = f_1(\mathbf{r} - \mathbf{m}) = f_1(\mathbf{r} - \tilde{\mathbf{m}}) + \nabla f_1(\mathbf{r} - \tilde{\mathbf{m}})\delta\mathbf{m}$$

Similarly, $\delta\mathbf{m}^{-1} = \tilde{\mathbf{m}}^{-1} - \mathbf{m}^{-1}$ and

$$\begin{aligned} \hat{f}_2(\mathbf{r}) &= f_2(\mathbf{r} - \mathbf{m}^{-1}) \\ &= f_2(\mathbf{r} - \tilde{\mathbf{m}}^{-1}) + \nabla f_2(\mathbf{r} - \tilde{\mathbf{m}}^{-1})\delta\mathbf{m}^{-1} \end{aligned}$$

Substituting this relation into the error equations, the following Euler-Lagrange equations for $\delta\mathbf{m}$ can be derived:

$$\begin{aligned} \lambda_I[\hat{f}_{1x}(f_2 - \hat{f}_1 + \nabla\hat{f}_1\delta\mathbf{m}) + \hat{f}_{2x}(f_1 - \hat{f}_2 + \nabla\hat{f}_2\delta\mathbf{m}^{-1})] \\ + \lambda_P(u - \underline{u} + \delta u) = \lambda_S(\nabla^2\delta u + \nabla^2\tilde{u}) \\ + \lambda_D(\delta u_{xx} + \delta v_{xy} + \delta w_{xz} + \tilde{u}_{xx} + \tilde{v}_{xy} + \tilde{w}_{xz}) \end{aligned}$$

$$\begin{aligned} \lambda_I[\hat{f}_{1y}(f_2 - \hat{f}_1 + \nabla\hat{f}_1\delta\mathbf{m}) + \hat{f}_{2y}(f_1 - \hat{f}_2 + \nabla\hat{f}_2\delta\mathbf{m}^{-1})] \\ + \lambda_P(v - \underline{v} + \delta v) = \lambda_S(\nabla^2\delta v + \nabla^2\tilde{v}) \\ + \lambda_D(\delta u_{xy} + \delta v_{yy} + \delta w_{yz} + \tilde{u}_{xy} + \tilde{v}_{yy} + \tilde{w}_{yz}) \end{aligned}$$

$$\begin{aligned} \lambda_I[\hat{f}_{1z}(f_2 - \hat{f}_1 + \nabla\hat{f}_1\delta\mathbf{m}) + \hat{f}_{2z}(f_1 - \hat{f}_2 + \nabla\hat{f}_2\delta\mathbf{m}^{-1})] \\ + \lambda_P(w - \underline{w} + \delta w) = \lambda_S(\nabla^2\delta w + \nabla^2\tilde{w}) \\ + \lambda_D(\delta u_{xz} + \delta v_{yz} + \delta w_{zz} + \tilde{u}_{xz} + \tilde{v}_{yz} + \tilde{w}_{zz}) \end{aligned}$$

These equations are solved via standard finite difference techniques using a steepest descent algorithm. At each step, $\hat{f}_1(\mathbf{r})$ and $\hat{f}_2(\mathbf{r})$ are recomputed and the best $\delta\mathbf{m}$ minimizing the weighted error terms is calculated. This motion field increment is added to the overall motion field and the procedure is repeated until the algorithm con-

verges to a solution.

Though the motion vector field describes a one-to-one mapping in a continuous domain, some subtleties are involved in the calculation of $\hat{f}_1(\mathbf{r})$ and $\hat{f}_2(\mathbf{r})$ on a discrete grid. The arguments, \mathbf{r} and \mathbf{m} , are defined on a uniform grid in the forward mapping of f_1 to obtain $\hat{f}_1(\mathbf{r}) = f_1(\mathbf{r} - \mathbf{m})$, but their sum is a real-valued vector. Therefore, the value of $f_1(\mathbf{r} - \mathbf{m})$ is obtained by trilinearly sampling values from the f_1 volume at the real-valued spatial location, $\mathbf{r} - \mathbf{m}$. Obtaining $\hat{f}_2(\mathbf{r}) = f_2(\mathbf{r} - \mathbf{m}^{-1})$ is a bit more difficult. Displacement of points in the f_2 volume effectively results in a cloud of unevenly spaced point samples of a continuous volume. Simply distributing these values into the nearest voxel on a uniform grid results in substantial artifacts in the resulting deformed dataset. To avoid these problems, a backward projection of the volume is used. For each point in the deformed volume, $\hat{f}_2(\mathbf{r})$, the immediate neighborhood of displaced voxels is sampled. The displaced voxels within this neighborhood are weighted by a Gaussian function according to the distance from the sampled voxel. The size of the neighborhood of displaced voxels is determined by the maximum magnitude of the flow vector field. To speed calculation, the size of the neighborhood is determined by computing the maximum flow vector field magnitude over subportions of the total volume. This limits the size of the sampling neighborhood in regions where the motion is small.

The prediction flow field $\hat{\mathbf{m}}(\mathbf{r}) = (\hat{u}(\mathbf{r}), \hat{v}(\mathbf{r}), \hat{w}(\mathbf{r}))$ was obtained using motion calculations from adjacent time segments. It is assumed that a particle in a volume between times t_i and t_{i+1} will move at approximately the same velocity between the volumes acquired at times t_{i+1} and t_{i+2} . A flow field approximating this motion can be obtained by deforming itself in a manner similar to the $\hat{f}_1(\mathbf{r})$ deformation. That is, we obtain the flow field estimate as $\hat{u}(\mathbf{r}) = u(\mathbf{r} - \mathbf{m})$, $\hat{v}(\mathbf{r}) = v(\mathbf{r} - \mathbf{m})$, $\hat{w}(\mathbf{r}) = w(\mathbf{r} - \mathbf{m})$. Likewise, a similar relation can be found to obtain a backward estimate of the flow field between times t_{i+1} and t_{i+2} based on the calculated flow field between t_{i+2} and t_{i+3} .

In addition to including the prediction term, it was found that because of the noisy characteristics of the individual PET gates and because of fairly large displacements with respect to the voxel size, a multiscale approach was useful to obtain suitable convergence. The $128 \times 128 \times 47$ volume was subsampled into $64 \times 64 \times 23$, $32 \times 32 \times 23$ and $16 \times 16 \times 23$ datasets using a uniform cubic B-spline approximation to a Gaussian pyramid [13]. A motion flow field solution was found at the lowest resolution, then was propagated at the next level as the initial condition of the flow field. This technique not only speeded the overall convergence, but in many cases it was found necessary to avoid solutions at incorrect local min-

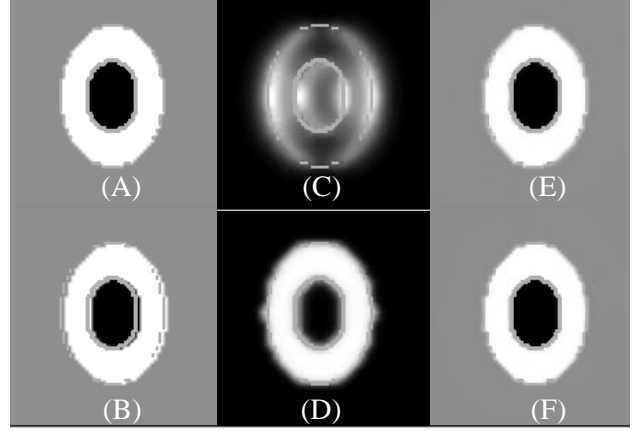


Figure 2. Translating Ellipsoidal Volume. Reference volumes (a,b), estimated flow fields for two cases (c,d) and resulting deformed volumes (e,f).

ima.

Once the deformed volume, $\hat{f}_1(\mathbf{r})$ matching $f_2(\mathbf{r})$ is obtained, subsequent processing to obtain a composite PET dataset is straightforward. The composite sum is computed as

$$f_{sum}(\mathbf{r}) = \hat{f}_1(\mathbf{r}) + f_2(\mathbf{r})$$

Because the deformed volume conserves the total counts present in the original volume (except at the volume borders), the composite volume represents the total PET counts acquired in the two gates. In general, data from all cardiac gates could be combined to form a single composite image. This paper will only consider the summing of two frames.

5. RESULTS

5.1. Simple Translation

As a demonstration of the technique on a simple object, the algorithm was applied to the translating ellipsoid on a noise-free background seen in Figure 2a,b. Figures 2c,d show the resulting x-component of the estimated flow field for two cases, the former using a small smoothness constraint weighting, the latter, using a large smoothness weighting. For the first case, where the smoothness constraint is small, it is seen that the motion estimate is nonzero only in the vicinity of the primary feature, i.e. the edges of the central ellipsoid. However, in the case using a large smoothness weighting, the flow field propagates so that the entire field is approximately a uniform translation. The deformed volumes, $\hat{f}_1(\mathbf{r})$, for each case are shown in Figures 2e,f. We see that in this example, the resulting deformed volumes for the two cases are similar even though the flow fields are markedly different. Obviously, this is true here because the background contributes noth-

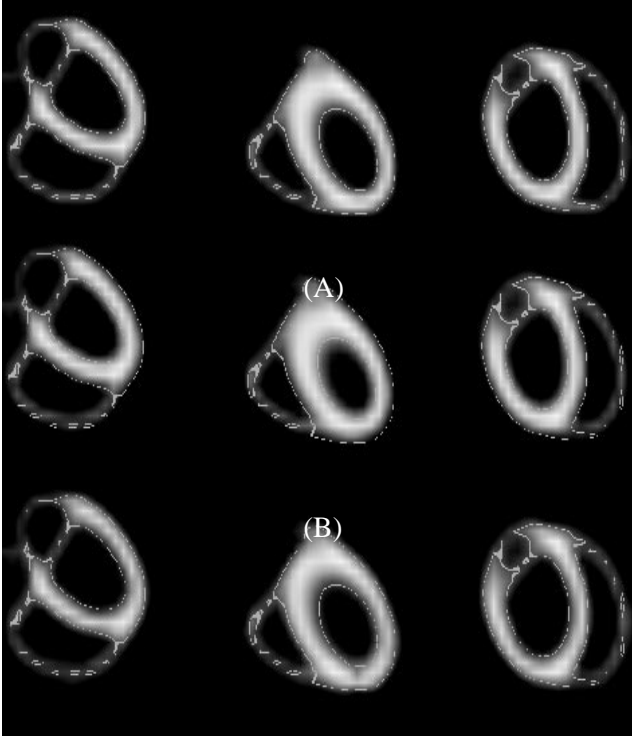


Figure 3. Three orthogonal slices of the MCAT cardiac phantom. Time frame f_1 is displayed in the top row, the next frame, f_2 , is shown in the middle row (A). Note the thickening of the ventricles and the inward contraction. The deformed volume of f_2 is seen in the bottom row (B) showing a good match with the target frame, f_1 .

ing to the deformed image. However, this characteristic is also partially true for the case of the “bright” ventricle on a dark noisy background. Accurate estimates of a deformed ventricular motion could be obtained using a large smoothness constraint that would enforce a smooth motion map even though the ventricular motion may be disjoint from adjacent material in the background, such as the blood pool and lung tissue.

5. 2. Simulated Cardiac Phantom

A realistic cardiac PET phantom was next used to test the algorithm. The phantom was obtained using the Mathematical CARDiac Torso (MCAT) software implemented at the University of North Carolina [14], [15], and provides a realistic dataset of the emission PET image of a beating heart against a zero background. Figure 3a,b shows three orthogonal slices through the dataset at two successive frames. Because of the noise-free nature of the data and the relatively small displacement with respect to the voxel size, a multi-scale approach was not necessary here. As a test of the 4D algorithm, flow estimates were made for four successive frames and were used to calculate a pre-

diction motion field for the gates displayed in this figure. The results, shown via the deformed volume, \hat{f}_2 , in Figure 3c, indicate that the algorithm appears to correctly track motion for this non-rigid case.

5. 3. Data From Canine PET Study

Two gated images acquired from the described dog study were selected for testing the algorithm. A gate during diastole was chosen as f_1 ; a gate at systole was chosen as f_2 . Three orthogonal slices through each dataset are seen in Figure 3. The motion estimation technique using only two time frames was carried out to obtain the appropriate mapping which warps the systole dataset to best match the volume at diastole. Figure 4 displays slices from the two gates, and the warped datasets, \hat{f}_1 and \hat{f}_2 . For these data a multi-scale approach was used, but the data were statistically good enough to obtain reasonable motion estimates without using prediction information from adjacent time frames. The deformed volumes in Fig-

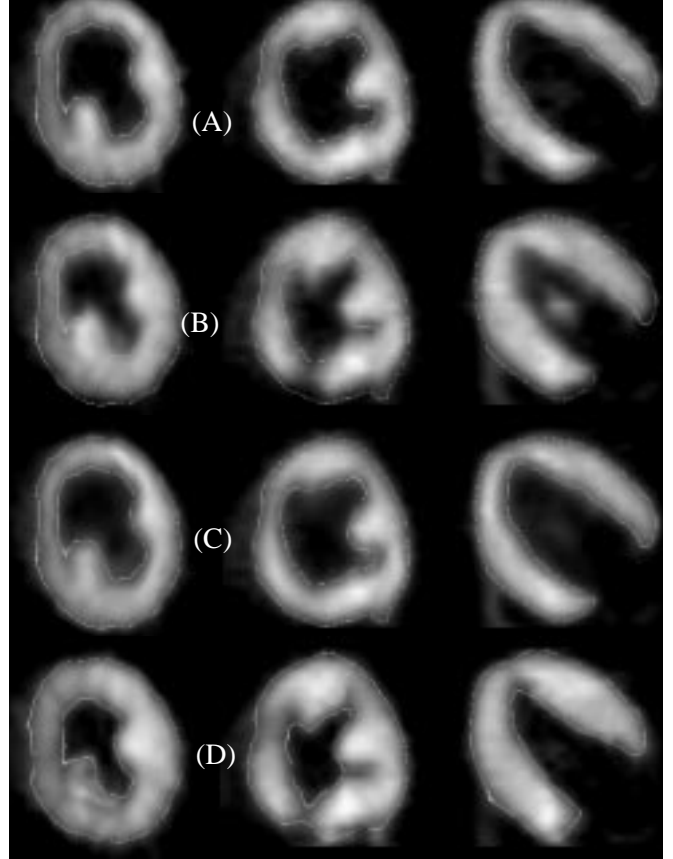


Figure 4. Three orthogonal slices of a gated cardiac canine study. A motion field is computed to register the diastole dataset (A) with data acquired near systole (B). The backward deformation of (B) is seen in the third row (C), and the forward deformation of diastole is seen in the bottom row (D). Both deformations provide a good match to the target volume.

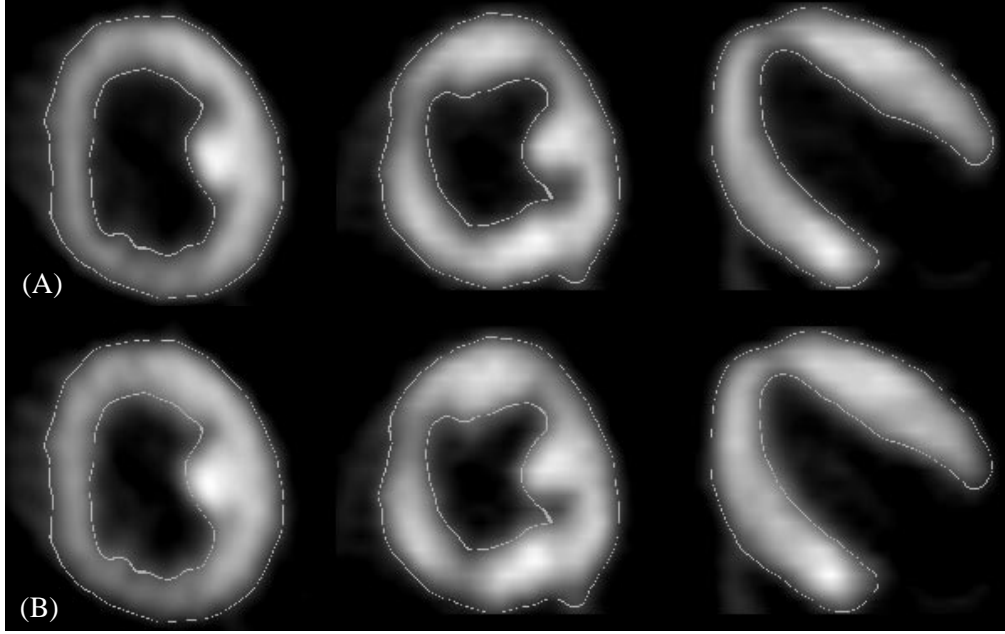


Figure 5. Combined gated data from cardiac dog study. Top row show images combined using motion compensation. In the bottom row, images are combined without motion compensation, inducing a blur proportional to the cardiac movement.

ures 4c,d show that a reasonable motion estimate was made to produce a dataset matching the heart shape at end diastole.

Results of combining the motion-corrected data are presented in Figure 5. The top row of images represent the data summed without prior motion compensation. Comparing with the same slices depicted in Figure 4a,b, it is

obvious that the contractile motion of the heart is blurred by such summing. The top row of images (Figure 5a) show the resulting sum using the motion corrected systole dataset. Motion-induced blurring is reduced.

5. 4. Data From Human PET Study

Transverse slices from two successive gates acquired

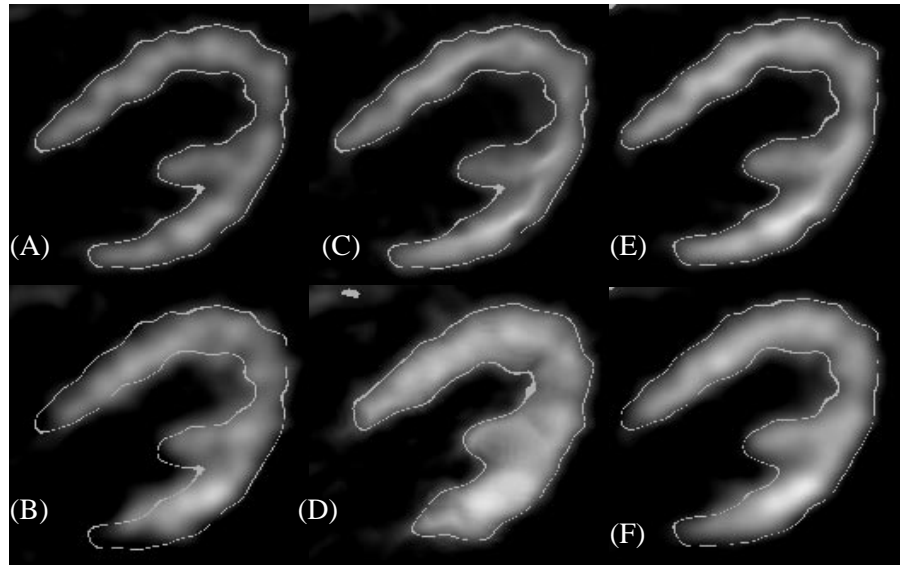


Figure 6. Human cardiac data (transverse slice) from two successive gates. Original frames f1 (A) and f2 (B) are used to calculate a motion field. Forward deformation of (A) results in image (D); backwards deformation of image (B) results in image (C). Adding original frame (A) with deformed image (C) results in the motion-compensated composite image (E). Less blurring is present than the non-compensated sum in (F).

from a human cardiac study are shown in Figure 6a,b. Using the notation defined previously, Figure 6a corresponds to f_1 , Figure 6b corresponds to f_2 . Deformation of the heart as it progresses from diastole to systole is evident by observing the edge mask of f_1 overlaid on f_2 . Because the individual gates of the human study were noisier than for the canine study, 4D priors were used in this solution. Estimates of flow fields for four successive gates were obtained using the multiscale techniques previously described. Once initial estimates were obtained, they were used to provide 4D information as a prediction field for the central gates displayed here. The resulting deformed volumes, \hat{f}_1 , and \hat{f}_2 obtained after calculation of the flow field are seen in Figures 6c,d respectively. Note that even though considerable deformation was required, as shown in Figure 7, the match between the deformed images and their reference is quite good. Results of summing images are shown in Figures 6e,f. The bottom image (6f) shows the result of a simple sum of the two gates without motion compensation; the top image (6e) shows the sum, $f_{sum}(\mathbf{r}) = \hat{f}_1(\mathbf{r}) + \hat{f}_2(\mathbf{r})$. Comparison of the edge map of $f_{sum}(\mathbf{r})$ overlaid on the uncompensated sum shows how the motion compensation reduces blur.

6. DISCUSSION

The motivation for development of this algorithm grew out of past experience which showed that though cardiac gating can stop motion due to the beating heart in practice, it is frequently of little consequence in PET imaging because the resulting gates are individually of poor quality due to limited statistics. We have demonstrated that by using a deformable motion model, it is feasible to effectively combine image data obtained from

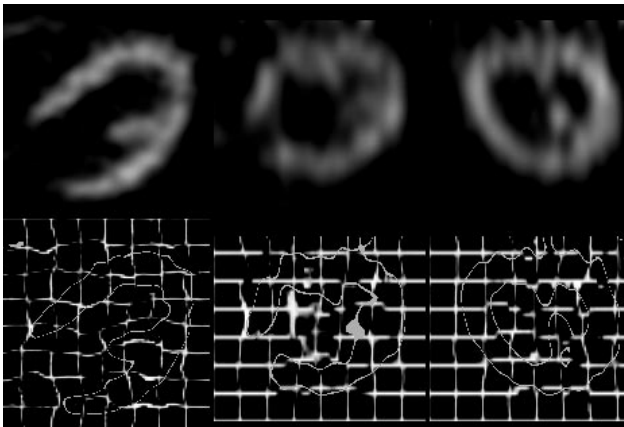


Figure 7. Motion field between two gates in a human cardiac study are displayed as the warp of a uniform grid. It is seen that considerable deformation is required to bring the images into register, yet the elastic constraints result in no topology changes.

different segments of the cardiac cycle. It is hoped that because of the improved statistics in the composite motion corrected image, quantitative measurements can be improved by this technique. We note that in its current form, the algorithm appears to only track the gross heart movement, primarily in a direction normal to the myocardial surface. It is known that complex motion occurs during the cardiac contraction. A nonuniform orientation of the myocardial fibers induces complex twists and stretches that are often tangential to the surface of the ventricle boundaries [16],[17]. Such motion would be hard to detect from PET data not only due to the spatial and temporal limitations of the detector, but also because of a lack of trackable features within the myocardium. Therefore, we do not propose that the motion model would be useful for strain and related kinematic calculations. However, when suitable features do exist, such as a local isotope uptake nonuniformity within the myocardium, it is expected that the algorithm will perform adequately because of the image matching and incompressibility constraints imposed.

As discussed in the background section, a number of researchers have already proposed algorithms for obtaining deformable motion estimates. However, because optical flow and other deformation techniques rely heavily on spatial and temporal derivatives, their success is questionable when applied to noisy datasets. This approach is unique in three separate aspects which were implemented to overcome the noisy characteristics of conventional gated cardiac data. First because the deformation operation intrinsically involves interpolation of offset voxel values, the forward and backward deformation will not produce identical error terms in the image matching constraint. For this reason, we simultaneously try to minimize the matching error of both the forward and backward transforms. The consequence of the procedure is a more robust estimation of the flow field that can be used to produce either deformed volume, namely, \hat{f}_1 , and \hat{f}_2 . Second, a multiscale approach is used to both speed the convergence of the algorithm, and to avoid local minima. Finally, the technique makes use of optical flow information from adjacent time slices to obtain a prediction flow field for use in the estimation.

Comparing the smoothness and incompressibility terms of our formulation with the elastic material model term of Bajcsy, we can note similarities. Both constrain the divergence and the smoothness of the motion flow field. At this point, we have not researched the most effective settings for the divergence and smoothness terms. Presumably, one would want to match the parameters to best characterize the elasticity of the myocardium. It should also be noted that this model assumes that the object in the field of view is homogeneous, and we know for a fact that

this is not true. For example, the elasticity of the blood pool and the lung tissue is considerably different from the myocardium. Additionally, there are areas where adjacent media are not connecting, and therefore a shearing motion would be possible. Fortunately for the case of PET data, these adjacent regions do not contribute significantly to the image data, so such issues are not a problem. For modalities where this characteristic would be true, such as MRI, one may have to employ nonisotropic elasticity models.

7. ACKNOWLEDGMENTS

This work was supported in part by the National Heart, Lung and Blood Institute of the U.S. Department of Health and Human Services under grant HL25840 and in part by the Director, Office of Energy Research, Office of Health and Environmental Research, Medical Applications and Biophysical Research Division of the U.S. Department of Energy under Contract no. DE-AC03-76SF00098.

8. REFERENCES

- [1] R. P. Woods, J. C. Mazziota, and S. R. Cherry. MR-PET registration with automated algorithm. *J Comput Assist Tomogr*, 17(4):536–546, 1993.
- [2] B. A. Ardekani, M. Braun, I. Kanno, B. F. Hutton, and I. Hidehiro. A fully automatic multimodality image registration algorithm. *J Comput Assist Tomogr*, 19(6):615–623, 1995.
- [3] B. Horn and B. G. Schunck. Determining optical flow. *Artificial Intell.*, 17:185–203, 1981.
- [4] J. M. Fitzpatrick. The existence of geometrical density-image transformations corresponding to object motion. *Comput. Vision, Graph., Image Proc.*, 44:155–174, 1988.
- [5] S. M. Song and R. M. Leahy. Computation of 3-D velocity fields from 3-D cine CT images of a human heart. *IEEE Trans Med Imag*, 10(1):295–306, 1991.
- [6] S. M. Song, R. M. Leahy, D. P. Boyd, B. H. Brundage, and S. Napel. Determining cardiac velocity fields and intraventricular pressure distribution from a sequence of Ultrafast CT cardiac images. *IEEE Trans Med Imag*, 13(2):386–397, 1994.
- [7] Z. Zhou, C. E. Synolakis, R. M. Leahy, and S. M. Song. Calculation of 3D internal displacement fields from 3D X-ray computer tomographic images. *Proc R Soc Lond A*, 449(1937):537–554, 1995.
- [8] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Comput. Vision, Graph., Image Proc.*, 46:1–21, 1989.
- [9] G. E. Christensen, R. D. Rabbitt, and M. I. Miller. 3d brain mapping using a deformable neuroanatomy. *Phys Med Biol*, 39:609–618, 1994.
- [10] G. E. Christensen, M. I. Miller, and M. W. Vannier. Individualizing neuroanatomical atlases using a massively parallel computer. *Computer*, 29(1):32–38, 1996.
- [11] P. Thompson and A. W. Toga. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans Med Imag*, 15(4):402–417, 1996.
- [12] C. Davatzikos. Spatial normalization of 3d brain images using deformable models. *J Comput Assist Tomogr*, 20(4):656–665, 1996.
- [13] B. W. Reutter, G. J. Klein, K. M. Brennan, and R. H. Huesman. Acquisition and automated 3-D segmentation of respiratory/cardiac-gated PET transmission images. In A. Del Guerra, editor, *1996 IEEE Nuclear Science Symposium Conference Record, Anaheim, CA November 2-9*, pages 1357–1361, 1996.
- [14] J. A. Terry, B. M. W. Tsui, J. R. Perry, J. L. Hendericks, and G. T. Gullberg. The design of a mathematical phantom of the upper human torso for use in 3D SPECT imaging research. In D. C. Mikulecky and A. M. Clarke, editors, *Biomedical Engineering: Opening New Doors (Proceedings of the 1990 Fall Meeting of the Biomedical Engineering Society, Blacksburg, VA)*, pages 185–190, New York, 1990. New York University Press.
- [15] J. A. Terry, B. M. W. Tsui, J. R. Perry, and G. T. Gullberg. A three-dimensional mathematical phantom of the human torso for use in SPECT imaging research studies. *J Nucl Med*, 31(5):868, 1990.
- [16] G. D. Meyer, M. C. Ziskin, W. P. Santamore, and A. A. Bove. Kinematics of the beating heart. *IEEE Trans Bio Engr*, 27(6):319–329, 1980.
- [17] A. A. Young and L. Axel. Three-dimensional motion and deformation of the heart wall: Estimation with spatial modulation of magnetization - a model-based approach. *Radiology*, 185:241–247, 1992.